

CLAIMS

1. A method for providing routing capability for variable-length unit traffic in a fixed-length cell routing switch with a shared memory

characterized in that

- an incoming variable-length unit, a frame, is segmented into a plurality of fixed-length cells including a start-of-frame cell, one or more continuation cell(s), and an end-of-frame cell,
- said fixed-length cells are routed through said switch,
- at an output of said switch, subsequent and deadlock-free transmission of consecutive cells of a certain frame is guaranteed by blocking any cell of a different frame from interleaving, which is governed by
 - defining a predetermined first condition for said shared memory,
 - defining a predetermined second condition for said certain frame and its cells,
 - controlling a cell's entry into said shared memory depending upon said first condition, and
 - controlling transmission of a said cell of said certain frame depending upon said second condition.

2. The method according to claim 1, wherein further

- a frame to be transmitted is assigned a multicast vector indicating a plurality of outputs of the switch.

3. The method according to claim 1 or 2, wherein further

- a frame to be transmitted is assigned a priority and,
- at an output of the switch, transmission of cells of a frame with a lower priority is interrupted when a frame with a higher priority is to be transmitted at said same output.

4. The method according to claim 3, wherein
- transmission of cells of the frame with the lower priority is resumed when the frame with the higher priority is fully transmitted.
5. The method according to any of the claims 1 to 4, wherein
- the first condition controls a cell's entry into the shared memory by granting entry if and only if
 - there is space available in said shared memory
 - AND
 - a predetermined "almost full" condition, indicating that a limited memory space is available, is not met in said shared memory OR an "active frame" condition is met, indicating that a certain frame's start-of-frame cell has been transmitted, but not its end-of-frame cell.
6. The method according to any of the preceding claims, wherein
- the second condition allows a cell's transmission through the switch to a certain output if and only if
 - the output queue associated to said certain output has not exceeded its pretermined threshold
 - OR
 - said cell is marked "active" for said certain output, i.e. belonging to an "active frame" whose start-of-frame cell has been transmitted, but not its end-of-frame cell.
7. The method according to claims 5 or 6, wherein each output
- keeps track which frame is currently "active", i.e. being transmitted, and
 - will not process any other frame until the end-of-frame cell of said current one is transmitted.

8. The method according to any of the preceding claims, wherein

- each cell is assigned a flag indicating whether it is a start-of-frame or continuation or end-of-frame cell.

9. The method according to claim 5, wherein

- the "almost full" condition for said shared memory is a programmable threshold below a "memory full" threshold, i.e. the total space available in said memory.

10. The method according to claim 9, wherein

- the "almost full" threshold T_{AF} is smaller or equal to the "memory full" threshold T_{SM} minus D times the number of ports N of the switch,

$$T_{AF} \leq T_{SM} - D \cdot N,$$

D being the flow control delay in cells in the switch.

11. The method according to any preceding claim, wherein

- the second condition controlling a cell's transmission through the switch to a certain output is extended to include multicast deadlock prevention by defining a frame to be "active" if and only if its last copy is being transmitted, so that said second condition allows a cell's transmission if and only if

- the output queue associated to said certain output has not exceeded its pretermined threshold

OR

- said cell is marked "active" for said certain output, i.e. belonging to the last copy of an "active frame" whose start-of-frame cell has been transmitted, but not its end-of-frame cell.

12. The method according to any preceding claim, wherein

- the second condition controlling a cell's transmission through the switch to a certain output is extended to include frame priority deadlock prevention by defining a frame to be "active"

according if and only if its highest priority is being transmitted, so that said second condition allows a cell's transmission if and only if

- the output queue associated to said certain output has not exceeded its pretermined threshold
- OR
- said cell is marked "active" for said certain output, i.e. belonging to the highest priority of an "active frame" whose start-of-frame cell has been transmitted, but not its end-of-frame cell.

13. A fixed-length cell routing switch providing routing capability for variable-length cell traffic with a shared memory,

characterized in that each address of said memory includes

- storage space for one cell, and
- an associated multicast counter,

thus, when a cell arrives, providing means for

- storing the cell's data at an available memory address,
- storing said used address in each destination output queue, and
- counting and storing the number of destinations in said counter.

14. The routing switch according to claim 13, further including

- a plurality of registers, each associated with a particular input and a particular output of said switch for storing an indication, preferably a bit, indicating that said particular output is "active" with respect to said particular input,
- said indication being set when a start-of-frame cell of a specific frame received on said particular input is transmitted on said particular output and said specific frame being the last copy to be transmitted, and
- said indication being reset when the corresponding end-of-frame cell is transmitted at said particular output.

15. The routing switch according to claim 13 or 14, further including

- means for assigning a priority to a frame to be transmitted and
- means for interrupting and/or resuming transmission of cells of a frame with a lower priority whenever a frame with a higher priority is to be transmitted at the same output.

16. The routing switch according to claim 13, further including

- a plurality of registers, each associated with a particular priority, a particular input, and a particular output of said switch for storing an indication, preferably a bit, indicating that said particular output is "active" with respect to said particular input and said particular priority,
 - said indication being set when a start-of-frame cell of a specific frame of said particular priority, received on said particular input, is transmitted on said particular output and said specific frame being the last copy to be transmitted, and
 - said indication being reset when the corresponding end-of-frame cell is transmitted at said particular output.
- a plurality of registers, each associated with a particular priority, a particular input, and a particular output of said switch for storing an indication, preferably a bit, indicating that said particular output is "transmit active" with respect to said particular input and said particular priority,
 - said indication being set when either a cell of a specific frame of said particular priority, received on said particular input, has been transmitted on said particular output in the current packet cycle, or no cell has been transmitted at all on said particular output, and
 - said indication being reset otherwise.